



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/981,620	10/16/2001	Richard L. Coulson	5038-118	6345

8791 7590 03/21/2007
BLAKELY SOKOLOFF TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1030

EXAMINER

VERBRUGGE, KEVIN

ART UNIT	PAPER NUMBER
----------	--------------

2189

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
2 MONTHS	03/21/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

MAILED

MAR 21 2007

Technology Center 2100

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 09/981,620
Filing Date: October 16, 2001
Appellant(s): COULSON, RICHARD L.

Matthew C. Fagan, Reg. No. 37,542
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 12/31/06 appealing from the Office action
mailed 7/17/06.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

No amendment after final has been filed.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

Art Unit: 2189

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

Cohn, et al. UK Patent Application GB 2,286,267 A. Sept. 8, 1995.

IBM Technical Disclosure Bulletin NN9411421. Nov. 1, 1994.

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

Claims 107, 109, 110, 115, 118, 122-126, 128, 129, 134, 137, 139, 140, 145, and 148 are rejected under 35 U.S.C. 102(b) as being anticipated by UK Patent Application GB 2,286,267 A to Cohn et al.

Regarding claims 107, 126, and 137, Cohn shows the claimed elements in Fig. 1.

He shows the claimed non-volatile cache memory as cache memory 106. He teaches that this cache is non-volatile at page 4, lines 21-22 and at page 5, line 29.

Further, Cohn shows the claimed controller as control system 108 and described as cache control or management logic 108 at page 5, lines 30-31. As taught by Cohn at page 5, line 31 and following, "The cache control system 108 includes a suitable microprocessor together with associated control code to enable it to perform the various functions described herein...." Those functions include the claimed spinning down of the disk (page 2, lines 1-2; page 14, lines 10-15), queuing operations for the disk while the disk is spun down (page 2, lines 13-19 and 25-26; page 11, lines 17-19 and line 34 through page 12, line 4; page 13, lines 1-2 and 31 through page 14, line 2), spinning up the disk to satisfy a read request in response to a read miss (page 12, lines 28-31), and performing queued operations in response to the miss as claimed (page 12, lines 31-32).

In these passages and others, Cohn teaches that to achieve one goal of his device (minimizing the number of disk activations, page 3, lines 4-5), queued data is transferred from the cache to the disk once the disk has been spun up to satisfy a required disk access (such as a read miss) as claimed.

Regarding claims 109, 128, and 139, Cohn's device spins down the disk after the read and the one or more queued operations are completed.

Regarding claims 110, 129, and 140, Cohn's queued operations include a write operation as mentioned at page 2, lines 13-14 (disk accesses which are caused by the need to maintain consistency between the data stored in the cache and on the disk are due to writes to the cache that need to be transferred to the disk to maintain consistency).

Regarding claims 115, 134, and 145, Cohn's device operates as claimed.

Regarding claim 118, Cohn explicitly teaches "The cache control system 108 includes a suitable microprocessor together with associated control code to enable it to perform the various functions described herein...." (emphasis added).

Regarding claim 122, Cohn's memory controller processes digital signals and is therefore a digital signal processor.

Regarding claim 123, Cohn's memory controller is an integrated circuit with a specific application (controlling memory) and is therefore an ASIC

Regarding claim 124, Cohn's controller 108 resides with the cache in cache system 102 and in data storage apparatus 110.

Regarding claim 125, Cohn's controller 108 is shown separate from the cache and the hard disk as claimed.

Regarding claim 148, Cohn does not limit his device to any of the claimed devices, rather he teaches that his device is useful wherever it is desired to save power with a rotating disk, which includes the devices claimed.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 111-114, 119-121, 130-133, and 141-144 are rejected under 35 U.S.C. 103(a) as being unpatentable over UK Patent Application GB 2 286 267 A to Cohn et al.

Regarding claims 111-114, 130-133, and 141-144, Cohn does not mention prefetching, but Official Notice is taken that prefetching was well-known in the art at the time of the invention. It would have been obvious to one of ordinary skill in the art at the time the invention was made to implement prefetching in Cohn's device to improve system operation by fetching data before it was needed to reduce operation time (the known benefit of prefetching).

One of the more effective uses of prefetching is for sequential streams. Once a processor determines that a request is part of a sequential stream, prefetching is implemented to obtain subsequent data of the sequential stream before it is actually needed so that when it is actually needed, it already resides in the cache and can be accessed quickly from the cache. If a request is not part of a sequential stream, prefetching may or may not be desirable (overly aggressive prefetching results in storing data in the cache that will never be used, forcing data in the cache that would have been used again to be thrown out). Prefetching is always a design tradeoff between gaining the speed advantage of having prefetched data in the cache before it is actually requested and throwing out data that will be used again to make room for prefetched data that might not be used. The small size of a cache is what makes prefetching potentially more detrimental than beneficial to operating speed.

Cohn does not teach determining if queued operations are desirable and then performing only the operations that are desirable. However, Official Notice is taken of queue operation techniques whereby more recent queue entries make older queue entries obsolete and therefore undesirable. Those undesirable queue entries are then deleted to avoid wasted operations. This typically includes memory requests to the same address where a first write to a certain address is made obsolete by a later write to the same address, for example. Since the first write is still in the queue (and has therefore not been written to memory) when the second write to the same address is placed in the queue, the first write can be deleted with no consequence to program operation as long as there are no intervening reads to that same address.

Regarding claims 119-121, Cohn may not explicitly disclose the claimed arrangements of hardware, software, and drivers, but these are mere implementation arrangements of the claimed operation and their particulars are matters of design choice. The skilled artisan possesses the knowledge required to implement Cohn's device and to modify it as his particular situation requires. Whether to implement certain features of Cohn in software or hardware is a matter of design choice.

Claims 116, 117, 135, 136, 146, and 147 are rejected under 35 U.S.C. 103(a) as being unpatentable over UK Patent Application GB 2 286 267 A to Cohn et al. in view of the IBM Technical Disclosure Bulletin NN9411421 published 11/1/94, hereinafter simply the TDB.

Cohn does not teach that his non-volatile memory is a polymer or ferroelectric memory, however it would have been obvious to one of ordinary skill in the art at the time the invention was made to make it so for the attendant advantages of polymer ferroelectric memory.

The TDB teaches that it was known to use polymer ferroelectric memories for nonvolatile storage purposes. As taught by the TDB, polymer ferroelectric memory was a known type of nonvolatile memory at the time of the invention and it therefore would

have been an obvious choice to use for the nonvolatile memory in Cohn's device. It would have been obvious to one of ordinary skill in the art at the time the invention was made to use a polymer ferroelectric memory for the design benefits that it provides, including small size, inexpensive construction, and fast response times.

(10) Response to Argument

On page 5, fourth full paragraph, Applicant states that the Examiner's final Office action "appears to equate Cohn's designating sectors as New in cache memory 106 with queuing operations for a rotating storage device (claims 107 and 126) and with queuing disk operations (claim 137) as claimed by Applicant." Applicant continues in the next paragraph, "A designated state of a sector in cache memory 106 cannot be equated with an operation. Simply put, a status is not an operation."

The passage of the final Office action cited by the Applicant is copied here (emphasis added here):

Cohn explicitly teaches at page 10, line 34, that destaging means "transferring data from the cache to disk." This is one of the two main tasks performed by the cache replacement mechanism in Cohn. The other main task performed by the cache replacement mechanism is replacement of items within the cache. This replacement is what anticipates the claimed queuing.

Replacement of items in the cache is the basic cache function of storing data at least until it can be written to disk (destaged). For example, when a processor creates a new piece of data, it writes the new data to the cache and the cache indicates that the data just stored is new (also called dirty in the art). This indicates to the cache controller that the data has not been written to the disk yet. This allows the processor to proceed with other operations and the cache controller writes (destages) the new data to the disk at a later time. Cohn describes the "new" designation at page 7, lines 11-12 and 20-23.

From the underlined passages, it is clear that the Examiner is not equating Cohn's designating of sectors as new with queuing operations. Rather the Examiner is equating the replacement or storing of items in the cache until they can be written to disk with the claimed queuing operations. Clearly a status is not an operation, and the Examiner had never asserted such. It is also clear that storing data in the cache when the disk is spun down and marking that data as "new" so it will be written to disk later when the disk is eventually spun up is a queuing operation and meets the language of the claims.

On page 6, first full paragraph, Applicant argues that "arranging of sectors in a Destaging Order as taught by Cohn cannot be equated with queuing any operations." The Examiner agrees, but the Examiner had never asserted otherwise. It was and is the Examiner's position that storing data in the cache when the disk is spun down and marking that data as "new" so it will be written to disk later when the disk is eventually spun up is a queuing operation and meets the language of the claims. Cohn doesn't explicitly use the word "queuing" for this storing of data in the cache, but Cohn's disclosure meets the broadly claimed function of queuing since queuing just means holding something until later and Cohn's device clearly holds the data in the cache until it can be written to the disk later when the disk is spun up.

Finally, in regard to section 1 of Applicant's arguments, on page 5, last full sentence, Applicant submits that "the term queue and its derivatives connote order." This single sentence is apparently an attempt to clarify the meaning of the claim terms "queue" and "queuing." The specification gives no clear definition of the terms except to

say that “queued memory operations may typically include writes to the disk and prefetch read operations from the disk” (page 4, line 25 through page 5, line 1 and page 7, lines 20-21). While a “queue” is mentioned in the specification, it is not shown in the figures, and one is left to presume it is stored in non-volatile cache memory 14 since it is clearly not stored on disk memory 12 and there is no other storage means shown in Fig. 1. As such, this “queue” is not well defined and in fact appears to be the cache or part of the cache. In any case, the claims do not require any special structure or location for the queue. All that is required in the independent claims is “queued operations.” Since the specification only mentions writes and prefetch reads as the examples of queued operations, these are presumably what is being claimed in the independent claims (writes and prefetches are specifically called out in dependent claims). As such, all that is required to meet the independent claims is a write operation, and clearly Cohn teaches as much (storing data in his cache and marking it “new” is clearly a write operation).

In section 2 of Applicant’s arguments, Applicant merely asserts that Cohn does not teach the claimed features and asserts that the Examiner’s final Office action “does not address any queuing of a prefetch operation or any performing of a queued prefetch operation in response to a miss.”

In response, it is noted that the Examiner explicitly stated that “it would have been obvious . . . to implement prefetching in Cohn’s device to improve system operation by fetching data before it was needed to reduce operation time (the known benefit of prefetching).”

Since Cohn's device includes queuing of memory operations (namely write operations) and performing queued operations in response to a miss (once the disk is spun up to satisfy a miss, the queued operations are destaged or written to the disk), and since it would have been obvious to implement prefetching in Cohn's device, it would also have been obvious to queue the prefetch operations and to perform queued operations in response to a miss. Applicant presents no evidence whatsoever why this wouldn't be obvious.

In fact, since it would have been obvious to include prefetching and since Cohn's device already performs queuing, it would have been obvious to queue the prefetch operations as well as the write operations. Prefetches are early requests for data that isn't needed yet and therefore prefetches are "low priority" or "low need" requests. The lack of need for a prefetch to be satisfied immediately means the disk would not need to be spun up immediately after a prefetch is requested. Cohn teaches that to save power, the only reason his disk is spun up is to satisfy a miss -- a request that cannot be satisfied from the cache. Therefore, since prefetches have such low priority, they can be queued and satisfied or performed later when the disk is spun up for satisfying a miss. So it would have been obvious to one of ordinary skill in the art at the time the invention was made to queue up the prefetch operations and perform them later when the disk is spun up to minimize the number of disk activations and thereby save the most amount of power (Cohn's chief goal).

In section 3 of Applicant's arguments, Applicant presents no arguments whatsoever, relying instead on the independent claims and their arguments.


Art Unit: 2189

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.


For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,




Kevin Verbrugge
Primary Examiner
AU 2189

Conferees:



Donald Sparks
Supervisory Patent Examiner
AU 2187



Reginald Bragdon
Supervisory Patent Examiner
AU 2189

DONALD SPARKS
SUPERVISORY PATENT EXAMINER

REGINALD BRAGDON
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100